Transcript

Moderator: Dr. Andy Packham

0:05
Welcome to Elevate.
Today, we're going to be talking about what it takes to move AI from pilot into production and creating impact. And that's real impact. Many enterprises have got many, many proof of concepts.

Some are active, some are stalled and some have got into production.

The real challenge is not about setting up the proof of concept, is not about ideas or use cases.
It's about operationalising AI safely, repeatedly and at scale.

Ultimately, it's about driving real business impact, boosting efficiency, enhancing decision making, unlocking measurable business ROI.

These are measures that your CEO recognises as having a real impact on the business.

So today we're going to be talking about how Microsoft Azure and HCLTech’s AI Foundry, turning those challenges into board level outcomes and how with the right governance, data, foundations and agentic passion patterns, you can deliver measurable business impact.
We've seen real success.
We've seen many of our customers really driving significant meaningful change.
Where they've been successful is where they have focused on driving change rather than in implementing technology.
And the real differentiator, I think is focusing on changing the process, the culture and the ways of working.
So today I'm really happy.
I'm joined again by Jeeva AKR from Microsoft and Srini Kompella from HCLTech.
Not only are they returning guests, they were, we were talking 18 months ago, but also they're both industry experts.
They were both worked consistently to cut through the hype and focusing on driving real world insights and practical guidance.
So yeah, like I said, it's 18 months.
So Srini, maybe starting with you, 18 months is a long time.
What have you seen change?

Speaker: Srini Kompella

2:13
Yeah, 18 months is, you know, is almost like an era now.
I mean, 18 months ago.
I mean, there's still a lot of conversation about generative AI, you know, the POC’s, the pilots, how, how to make it work.
Now we are, you know, talking about Agentic AI.
But overall it's, it's really about, you know, how do, how does AI get implemented at an enterprise scale to drive the value?
I mean, that's really where I think the, you know, the companies are beginning to shift.

It's no longer about, you know, the small POC's and pilots.

It's no longer about the five-minute efficiencies.

You know what I mean by five minute efficiency is that, you know, you use an AI piece of an AI technology and you do something a little bit faster that saves you five minutes.

But that doesn't necessarily, you know, not always, you know, converts into business value.
So I think there is a lot of the growing realization that, you know, companies need to think about enterprise, you know, level AI.

It's not about the small POC’s And as a result of that, there is a growing realization now that, you know, in order to do that enterprise level AI, there are a lot of things that need to come in place, as you rightly said, Andy, I mean, it's, you know, it's change, it is rewiring some of the business processes.

It is rewiring, you know, several of the IT processes. And if you want to derive, you know, large scale business value from AI, you know, the thing that powers that AI is data together with insights, right? So either you create your insights on your own with AI or you got to create the insights for AI to be, you know, impactful.

So there is this realization coming in that, you know, companies need to take a holistic view to the, you know, infrastructure, to the data estate, to the analytics estate in order to power AI, you know, and, and put it in the hands of almost every single employee in a company.

That's the big shift, you know, that we are seeing of course Agentic AI, you know, is, you know, potentially accelerating it.

But without that foundation, you know, it is very hard for companies to, you know, deliver consistent repeatable business value with AI.

Moderator: Dr. Andy Packham

4:21
Yeah, that's acutally.
Again, Jeeva what have you seen in 18 months?
It's a long, long time in the AI world, isn't it?

Speaker: Jeeva AKR

4:29
Yeah.
First of all, thank you so much for having me, Andy and Srini, wonderful to work with both of you and partnering with you in this conversation.
So I will first of all, great question.
18 months is absolutely a long time in an industry that has got no franchise value, right?
Things change at a rapid scale, especially in the age of AI these days.
So I will probably frame this conversation in in three different ways about the what's we are actually

what we are observing as a change.
One, if you think about it about 18 months ago majority of our customers were actually moving from transitioning from RPA based technologies to single Agentic AI systems.
And, and we all know about the single Agentic AI systems, which is about more focused on content search and content summarization, content creation and all that.

These were the predominant use cases that people were actually trying to use for now what we are seeing is that I think Srini used the word Agentic AI systems.

We have seen the evolutionary process in terms of organization, how organizations are thinking about agents being multi Agentic AI systems, which means that multiple agents working autonomously, independently, collaboratively to solve the complex problems, right.

I think that's a a big state that we are actually venturing into.
So that's from the overall industry standpoint, that is a big shift in the way that the thought process has gone about.

And then the second thing is if you think about how customers actually went about AI 18 months ago or even until recently, majority of the conversation anytime when we have gone into customers, it was always focused on the choice of model, the model availability, the cost of models and the API's is GPU availability, the cost of GPU.
These were the predominant conversations that were all centered around it.
And that's primarily because most of these companies were focused more on proving the value or POC at that point in time pilot, right.
So that's they were super tactically focused on just that one layer.
And now what we have seen is that there is a realisation like what Srini said, it's that it's not just the AI part.
It is it's about how you get your data foundation right, if you want to create a sustainable and scalable AI in the organization.
And then the second thing is about what's about your AI model, which is what you actually said it every, which is about how do you govern all the AI models?
How do you make sure that you are actually creating responsible AI use cases within your organization?
And then the third, don't leave the people behind because the best people who can actually work on your organization is about getting your, the entire team, whether it is a business user or people who are actually working on your AI project, it's about up levelling all of them towards your initiative.
So there is a realization for customers about how they are thinking about that.

So all of it.

Moderator: Dr. Andy Packham

Surely then, I mean, look, we talk a lot about Scale AI.
What, what does scaling, you know, what does scale AI mean though?
So what we mean by scaling AI is really, you know, AI at an Enterprise level.
So you know, it, it is about having AI in the hands of almost every employee in the organization.
Because end of the day, you know, AI is going to, I mean, Jeeva talked about it, right?
You know, uplifting the people, AI can actually uplift people, you know, much, much faster in, in executing their priorities, not just on a day-to-day basis, you know, and also in terms of, you know, transforming the business processes, right.
So you know, enterprise, I mean, when we say scaling AI, it is enterprise scaling of AI because if you, you take a couple of steps back and look at AI is going to be everywhere.

It's going to be in our, you know, in our handheld devices as consumers of AI, right?

Is, you know, look at it, you know, that category, you can call it as consumer AI, right?
What we're really talking about here is how can enterprises benefit with the AI?
So while the consumers are using more and more AI in their day-to-day lives on the, you know, phones, laptops, devices, you know, how does it, you know, a company scale AI at that, you know, enterprise level.
So first is people, how do you get AI in the hands of people, right, in order for them to uplift themselves, right?
Second is in terms of, you know, how do you set up an overall AI Estate, which then needs to include the Data Estate and the infrastructure so that you do, you know, the quote unquote use cases.

You know, you know, when I say quote, UN quote, it is because it's not about the small 5 minute efficiencies alone. It is about, you know, how do you do demand forecasting better?
How do you do your risk modelling better? How do you do your financial, you know, forecasting better and hundreds of use cases, you know, at scale, I mean, you know, without worrying about what is the ROI for each use case, right?

What amount of GPU goes into this particular use case versus that use case? Which model should I use for this use case versus that? Because that actually, you know, becomes a blocker for not just executing the projects, but also in terms of, you know, scaling of AI. So, you know, so the second part of it is that, you know, what kind of an integrated AI data estate that companies need to deploy so that you do these use cases at scale.

The third is, you know, the last mile integration. So, you know, it's more about, you know, creating an AI model. It needs to be integrated into the business processes, into the systems of record of systems of engagement in order for the people to derive better value with AI, right. So, all of these put together is what we call as the enterprise scaling of AI.

Moderator: Dr. Andy Packham

10:29
Of course.

Speaker: Jeeva AKR

I would like to, yeah, I would like to add to what Srini said. I want to maybe provide a little more context to what we are seeing.

OK.

You asked the question about what is scale AI and then we talked about how the customers are transitioning from their pilot POC's to full scale implementations and we talked about the central theme as a data foundation. So let me explain a scenario so that the listeners of this podcast will be will be able to get behind our line of thinking about what we are trying to get this to.

One: If you think about POC’s, that is anytime when customers do POC and pilot, it is very specific use case built on a very highly curated data sets that's sitting in a very confined environment.

Almost every single use case that you develop is going to give you ultra success on that one in that pilot environment because you're training the model on a highly curated data set in a confined environment. Now I'm just going to use an example of a financial industry. So most of ours, we know although we are not in the financial industry, all of us are financial services customers for one or the other financial institutions, right. And we have an understanding about all these financial institutions.

They have corporate banking, they have retail banking services, they have investment banking, they have customer service, all of these things, different banking services that they provide.

Now the when we go and deploy all these agents, the reality of the situation is each and every one of these set of financial systems which is corporate banking, retail banking and also the investment banking, they are being supported by nearly dozen to two dozen applications at any given moment for depending on the size of the organization.

Guess what, all these applications produce large amount of data that sits in either in SQL servers, Postgres, Mongo DB or you name it, any kind of data store, right?

And the challenge for all these customers is when you go on deploy all these highly trained data set, oh sorry agent which is on a highly curated data set and then you go and deploy it in a real world environment.

The real world environment is not highly curated because agent has to work across data that came from about nearly half a dozen to for two dozen applications depending on the size of the organization.

Now it needs to work independently and collaboratively. Think about the challenge that you have in there, right?
So it's the data proliferation that happens. That's the data silos. That's a major issue when we talk about all the all the agent AKI systems being deployed at scale.

So when we talk about Scale AI, yes, fundamentally we need to think about how we can empower agents to not just being trained on one highly called, highly qualified data set, but it is also about how you can actually help them to make decisions on a near quality, high quality data set, right.

So when we talk about high quality data set, it's the data foundation should be able to support open format, open standards. So that agent had doesn't have to do the all the data transformation when it comes from data comes from multiple different systems and multiple from applications, right.

So for example, customers use case might require the agent to leverage the application data that comes through the SQL Server, but also the click stream data or other log data that is coming from somewhere else from the apps too.
So when you're trying to combine it, the data, different data formats and all this actually produces a lot of challenge for the systems to scale.
So when we think about scale, the majority of the thing is first, how do you create a data foundation that solves for all this efficiency issues in terms of data transformation, keeping it all based on open standards and open, open format.

And then of course, Srini already talked about don't leave the people behind because you need to think about the operating model, the Agentic AI systems.

What majority of our customers understand now is that it is an iterative journey.

It's not that you can actually go and train an AI model and deploy it and the entire organizational challenge will get solved in one go, right?

It is an iterative journey, which means that first year every organization needs to keep a goal of what is the performance efficiency that they are actually looking to achieve.

What is the business outcome that they are looking to get through that Agentic AI deployment and having a measurable goal and a practical goal and then continuously monitoring it and measuring it for the business value and then iterating over it to increase it over the period of time.

That's provides a much more Scalable AI for organizations.

Speaker: Srini Kompella

15:32
And yeah, Yep, if I can just add to that a little bit, I think so.
I think that agentic AI, you know, proliferation, right?
Just want to add one more thing.
I mean, what we typically when we're talking about building or Scaling AI, we talk about to the point of, you know, developing and integrating AI, let us say with systems of record applications, business processes. The hard work actually starts after we put into the production as well.

So just, you know, want to kind of extend a little bit on what Jibo talked about, right?
In the real world, data keeps changing.
So you know, so there is, there's a data drift that keeps happening.
So which impacts the model, the model can drift as well.
So, you know, both of these need to be, you know, handled closely.
There could be, you know, changes in regulation, for example, if we are dealing with a regulatory environment, the changes in regulation will impact the data sets on the models, you know, as well.
So those need to be tweaked. I mean, you can call them as you know, model operations, but in the world of agent AI, things evolve a lot more.
So if you are talking about autonomous agents, right, so each individual agent, you can make them reliable, let us say to about 90-95%.
But if you want three agents to work together in sequence, just the way you know, 3 humans, you know, work together in sequence, if each agent has a reliability of, let us say 90%, the, you know, if you have three agents, the reliability of the chain is not 90%, it is about 74%.
So, you know, so which means that, you know, there's a continuous effort that needs to go in to enhance the reliability of the agents, right, for the humans to fully trust the, you know, Agent AI, right?

It is not just about AI agents, but the entire Agenttic AI system.

So when, you know, so now the, the, the big part of the scaling of AI is also in terms of how do, how do we make sure that AI reliability is increasing in production and in operations on a day-to-day basis, not just remaining static, right?
Increasing and able to handle the variability of the datasets.
You know, either because of changes that are happening internally within the company or because of external changes like the example of regulatory.
There's a big piece that is equally, if not more important for the enterprise scaling of AI.

Moderator: Dr. Andy Packham

17:51
Yeah, yeah, I think yeah, you're, you're, you're spot on.
It's, it's not easy, right?
It's, it's very complex and the world, you know, the world continuously changes.
So, yeah, Jeeva, but just can you just walk us through the Azure stack?
Well, you know, what are the core elements of the Azure Stack that we should be looking at in terms of building that reliability and cost control, obviously cost control into this?

Speaker: Jeeva AKR

18:14
Absolutely, Andy.
So first of all, I want to make it super simplified for everyone to understand.
So it's when we talk about stack, I don't want us to think about this being a very complex architecture with multiple things.
And in fact, any time when we go and talk to customers, one thing that we always say is that just the way we think about how we started this whole conversation, customers should not focus only on the AI part, which is just model, model availability, GPU availability, cost of GPU's and all that.
But it should be the underlying data foundation which solves for all the efficiency issues that we both talked about just now. And then on top of it, we already talked about the people, which means that, you know, there is a responsibility I part that comes in the skilling part of it that comes in and all that.
And then final piece is that the operating model which governs all of these elements, right?
So how do you, Srini talked about very nicely about one data quality.

Then the second is about your AI model accuracy, how you are actually progressing on it and how people are actually getting use, the use of AI and also the familiarity with all the tools and all that stuff, the data context, all of it like that, that's the operating model.

So in order to address all these challenges, yes, it's a complex problem.

And when, as I represent Microsoft, our intent is to think that we can solve everyone's problem through technology.
But the reality is that that is it's a Harvard study that I want to quote to everybody that you know, 90% of the problem related to transformation is on the people process culture within the organization. And only 10% can be solved with technology.

But within that 10%, we believe that we between 3 products, we can actually move the maturity curve for our customers in a big way.
And that is 1 fabric as a data foundation for majority of the majority of our customers because.
One of the key value proposition of Fabric is the One Lake we ask the data foundation.
Many customers misunderstand that One Lake is just another data storage.
It's just not.
It is actually a data foundation, which solves for a lot of the complexity the customer face when it comes to integrating all the data sources.
So we have naturally built the software layer to provide both the virtualization capability for our customers.

So customers who are actually having the data state that is spanning across multiple data systems, multiple geographies and multiple cloud environment.
Now we'll be able to have access to the data set without having to move out of the system from one lake. So that's a shortcut capability.

And then the second thing is the mirroring. And the mirroring is about, you know, the, the as we talked about all these applications produce a lot of data which is sitting on SQL Server, Postgres, Mongo DB, Cassandra, I mean all the data source that we talked about, you name any one of the database.

And the problem is that today the organisations simply rely extensively on ETL sources to pull all the data and bring it into a common for common data storage in order for them to get transformed and be exposed for other AI agents and applications.
But that kind of efficiency is not going to be working very well for AI agents, which are created to solve problems autonomously and based on the real time data that is coming into the system in near real time maybe.
So from that's the, for example, like the fraud detection, risk analysis and all that in a financial institution.

So we think about how do you actually create a mirroring capability where you bring in all the data naturally with the CDC capability so that you have all the data sets available for agents to work on top of One Lake Asset Data Foundation.

So I already talked about the fabric, the One Lake Asset Data Foundation at the bottom, the second or if you go on top of it, which is the AI foundry, which you and Srini you both, I referred it, which is about:

- How do you actually have the model choices?
- How do you the how can the customers can actually pick the right model for the use cases?
- How do you measure the performance of all the use cases?
- How do you measure the performance of all the agents that are actually getting deployed?
- How do you manage all the agents that are deployed across the enterprise, right.

So all of this requires an extensive observability capability beyond just the model catalogue and all that and AI Foundry solves for it. And many customers know already that AI Foundry is deeply integrated with the Fabric.

And then of course, on top of it, which is the copilot, which is I cannot underestimate importance of the value of the dare copilot that you have on top because.

For example, many of the customers always ask us about what is the difference between your M365 Copilot that Microsoft has and the Copilot Studio versus the one that I can get from ChatGPT OR from the outside world.

The thing is that the most powerful your data that you can have is the contextual data that you have within your enterprise, right?

So while data, you know, the, the AI models can go and source all the data that is outside in the that's present in the outside world.

But where Microsoft does a great job is about using your own enterprise data, the data that is sitting in your own systems like emails, the chat communications within your organization, within the SharePoint.

So combine that outside world with that inside knowledge that you already have, it becomes a very powerful proposition for customers to move forward anytime you do it.

So to answer your question about what is the Microsoft stack, it is the Fabric as a data foundation and then AI Foundry and Copilots running on top of it.

And of course, I already talked about the importance of Purview as a data governance layer. You cannot move forward without having a sound data governance layer that is governing all of these three. And one of the major challenge for customers, like I said, is, you know, they have a data state that's crawling across multiple data systems, multiple geographies.

How do you get a singular view of all the data asset? If you can't have a singular view, then you can't catalogue them.

And if you can't catalogue them, you cannot apply all the policies.
And if you can't apply policies, then you're not going to have a scalable AI in the safe AI within the organization. So all these four are very well integrated.

Andy, it's a long answer, but I hope that I provided the clarity and the context.

Moderator: Dr. Andy Packham

25:03
No, that too.

25:04
That's I think, yeah.
When you look at actually the complexity, I think that's a very, that's a very short answer, but thanks for that.

Srini we've built, you know, HCLTech's AI Foundry. I mean it's closely coupled into Azure, tightly built.

What makes, you know, what makes our AI Foundry different and helping customers move from pilot, you know, into a proper scale business.

Speaker: Srini Kompella

Yep.

25:29
So if we look at the HCLTech AI Foundry, I mean, it's got 3 distinct, you know, components on the value chain of, you know, data to insights to AI that's open by, you know, what we call as cognitive infrastructure, which is essentially, you know, a hybrid, you know, infrastructure, you know, platform

approach. Because we do believe that, you know, most companies will use a form of hybrid cloud or, you know, hybrid multi cloud or combined with, you know, some level of data sets and models sitting on-prem, right, for some specific, you know, scenarios.
So with that as the foundation, if we take the data, you know, analytics and the value chain, I mean, we do three things there.

One is autonomous data engineering.
I mean, Jeeva alluded to, you know, talked about, you know, data fabric, right?
And he alluded to the importance of, you know, getting all of the data, you know, you know, well put together in a simple way, non-technical way, just well put together that people can trust.
So that, you know, needs to have the right level of data quality, you know, right level of data governance so that people can trust the data.
So which is what we are now calling it as the autonomous data engineering.

It's really about how do you build the data engineering from all of the data sources, right? Whether it is SharePoint or ERP systems or, you know, other systems of record and pull that together into a data set that is, you know, more suitable for running your analytics and AI so that, you know, entire space. We could, you know, pull it together into autonomous data engineering.

What does it mean?
It means that, you know, data engineering is A: less people dependent and more focused on, you know, getting the data to be AI-ready or getting the data to be trusted.

That includes, you know, the usual ETL or ELT, you know, it includes, you know, data quality, it includes the verification, validation of the data because you know, you can, you're integrating various data sets, you got to verify and validate if these data sets are still, you know, accurate.

So all of it, we call it as autonomous data, right?

So that's number one.
And that then becomes the foundation for the second, which is, you know, the AI foundry augmented sites.
So today if you look at it, I mean, you know, the insights or analytics is an area which is, you know, quite complicated already, right?
Over the last decade, everybody has hundreds and thousands of reports, multiple technologies, you know, So as a result, which insight do you trust?
A simple thing as you know our pool, you know, in a telecom or sales reporting, if different people are building different insights without, you know, having a fully validated data sets or without having the right algorithms for the reports themselves, then you have two different insights.

So which one do you trust?
So, the second step really is in terms of, you know, how do you simplify the insights or augment the insights with AI so that, you know, reduces cost significantly, but more importantly, it increases the trust in the insights of the companies.
You know, one of the customers we are working with, for example, I mean, they have nearly 160,000 reports in just in the US.

It's a, it's a global multinational company, 160,000 reports.
You know, think about simply the ratio of reports to the, you know, employees of the users, right?
Do people really need so many reports, so many insights, which actually creates more confusion and chaos than actually helping?

So the second part of it is that, you know, how do you simplify insights or more critically, how do you augment insights with AI so that now you can focus on the Third aspect, which is the scaling of AI.

So in that within the foundry you got, we got two key components.

One is, you know, the agent hub.
It's really about how you build agents, you know, and orchestrate the interaction of the agents and also how do you, you know, address issues, potential issues.
I mean, not many customers are talking about it right now, but we kind of looking forward and saying how do you do BCP and DR with Agentic AI?

So, if an agent you know, does not perform to the right level compared to human, for example, let's say a human goes on a holiday or a vacation, there's a BCPDR plan.

So how do you do a BCPDR for agents, right?
How do you orchestrate the agents?

So, you know, within it HCLTech software?
I mean, we, we have, you know, a really cool product called UNO.
So which actually helps with orchestrating, you know, the interactions between the agents.
So we have the MIT number partnership, for example.

I mean, that again creates a different kind of options for, you know, agent interactions to happen with the different protocols.

So the first part of the scaling of the, you know, AI is the, you know, agent hub.

The second part is business persona-based agents.

So, so far, you know, we've been talking about, let us say agents in a little more of an abstract term, right?

I mean, what is an AI agent? What is an agent AI system?
At the end of the day, if the purpose of AI is to put AI in the hands of, you know, employees, which are mostly business users, right, then you know, we've got to create agents that are linked with business personas.
And that's the, you know, the second layer business persona based agents.
Now each agent, if you think about it, I mean, you know, we've, you know, that business persona based agents is classified into, you know, few categories, an **insight** agent, which helps a business user pull insights that are relevant for them, right?

A **knowledge** agent, an **action** agent and a **decision** agent.
So essentially, you know, it's an ecosystem of four agents that are mapped to a persona that enables a business user to do their day-to-day, you know, activities better, you know, with the agents.
So we are not yet looking at it as, you know, agents replacing humans, but it's really in terms of how agents can augment the business users to do their, you know, activities better, going beyond the five minute efficiencies and creating a business impact.

So these are the three distinct components of the, you know, HCLTech AI Foundry.

And if you think about it, I mean, it aligns quite nicely with what, you know, Jeeva was talking about as the, you know, the Microsoft strategy, right?

So it's a big synergy that we've got there in terms of autonomous data engineering, you know, augmenting insights with AI and then scaling of AI with a combination of agent hub and business persona based agents.

Speaker: Jeeva AKR

31:47
Beautiful.

Moderator: Dr. Andy Packham

31:50
I know we're kind of running out of time.
I wanted to try and squeeze a couple more questions if I could.
I'd like to get your kind of useful with all of this technology and experience.
Now, why do you think proof of concepts are, you know, are still stalling?
So I guess Srini, do you want to?

Speaker: Srini Kompella

32:09
Yeah.
So I think, you know, if if so, we've done our share of proof of concepts as a company, right?
You know, thousands of them, you know, or hundreds of customers across the globe.
So, you know, I mean, if you look at why are people doing proof of concepts, right?
So again, there's a difference between a proof of concept and proof of value.
So, but let's talk about proof of concept.
If we go back about 12-18 months ago, it was really about, you know, understanding the technology, right.
You know, what is AI? What is this new AI all about? So people have been trying to understand the technology itself.
Number one, second, understand how it can be applied right in in daily lives.
Third, you know, how do you actually operationalize it?
So, you know, a big chunk of POC's, you know, if, I go back and look at in year or so ago, we're more focused on really understanding the technology, how to adopt it, how to implement it, how to get value out of it.
So that experimentation has a cost, you know, and, and you know, I think the, one of the most recent studies has shown that 95%, you know, is for the positively 5% success rate, you know, with, with these experiments.
But there is probably a side benefit to it.
It's enabling people to understand what works, what wouldn't work, you know, what skills are needed for, you know, to really deploy AI.
So maybe, you know, when you look at it as a benefit of doubt on, you know, the POC, you know, situation in terms of percentage of success, right?
But that's primarily the reason people have been doing POCs, but that's shifting to proof of value.

So which is really saying that, you know, what value can we expect by applying AI, let us say into the KYC process, highly regulated, know your customer, right?
You know, so can you actually deploy Agentic AI to improve your KYC?

Can you, you know, deploy AI to improve your, you know, fraud detection, right?
Can you apply AI to improve demand forecasting?
Now, the proof of value means that, you know, there's a tangible business outcome that is expected.
So which means that, you know, there's a business case that need to be created now, right?
It's no longer about proving the technology to say that, hey, will this two work together and can I make it happen?
It's about, OK, if I put AI, you know, without using AI versus by using AI, can I see a tangible business value?
So it's, you know, I think there is still merit in doing the proof of value, not the proof of concept because you, you've got to assess, you know, what is the value by, you know, putting a certain technology in.
You wouldn't want to put a technology in just because it is there, it has to drive certain benefits.
So there's a definite, you know, benefit of, you know, continuing to do the proof of values, but in a more governed, in a more curated.
And most critically, I think Jeeva talked about it, you know, a few minutes ago.
It's not about, you know, some highly curated data sets. Can the proof of value be run on real world data?

Right.

So then things start becoming real then the chance of success, it's still a chance.
The chance of success of deriving the business value in production will go up significantly.

Moderator: Dr. Andy Packham

35:23
Jeeva, your thoughts?

Speaker: Jeeva AKR

35:25
Well, very well said.

35:26
Honestly, I think what Srini said is absolutely correct in terms of how we have seen the POCs and pilots.

If you think about it, like I said in the beginning, very specific data sets in a very highly confined environment will always produce magical results and during the POC and but I think that's a good suggestion coming from, you know, most of the customers have it in terms of maturity.
They train the model on a particular data set, but they always tested with the new different set of data sets other than the one that they trained on in order to go and validate all the models performance and all that stuff. But still it's not the real world data set.

So it's it will be a good suggestion, like what Srini said to go and see how you can bring a real world data set, maybe perhaps a week worth of data or a month worth of data and see how the models actually perform before actually going into the main school main scale production.
I think that's going to produce at least provide a better linear bridge between the theoretical world versus the practical world.

Moderator: Dr. Andy Packham

36:31
Yeah. So Jeeva, like kind of wrapping up now, right.

We've heard a lot of stuff. What advice would you give to anybody who's really wanting now to really accelerate that, that scale? What's the 123 things that maybe a customer should do if they really now they're going to really make a big push?

36:49
Absolutely.

Speaker: Jeeva AKR

36:50
So first, don't just focus on the technology, because technology like Srini said has been mostly solved for whether it is a tool availability, model availability and the GPU availability or the integration between multiple multitudes of tools.
I think majority of the problem related to that has been already solved for by a lot of few come few a few different companies, but it is the fundamental question about people process technology that you that you have within the organization, you need to create an operating model around it.

You need to think about how you can continue to measure the performance of each and every element that we just talked about, which is the data portion of the data quality, the data engineering, the along the data preparation time takes all of this should be a KPI and the data quality, that's the agents using data, their performance and efficacy of it that should be measured as well.

And then the final, the, the third is the, the business outcome value, which is about what is the, the value that you are actually getting out of all this Agentic AI systems.

And either in terms of the process efficiency that you are actually gaining or in terms of the experience that you are actually providing to your employees or to your users of the services or the products, right.
So there has to be a measurement and because unless you actually ground yourself on where you are right now, then you cannot actually continue to improve and see where the impact is going to be based on all the things that you are actually taking.

And then the last piece, which I have said it many, many times, don't leave your people behind because the best people who can work on the data are the pure people who have the context to the data.
Because many times AI fails because of the contextual information that is being missed in the data.
Because if you think about large organizations, they have data sitting from the last 30-40 years time.

Nobody knows at this point in time about which application produced to what kind of data and why they collected all the data and all that.

So if you don't have context and contextual information about that, then again, you're going to miss the boat in terms of the idea here that your AI is only going to be as good as your data, right.

Context matters and people are important part of the puzzle in that.
So that's my closing thought, Andy.

Moderator: Dr. Andy Packham

39:12
No, absolutely. Thanks.
Srini?
Yeah, same for you. Accelerate.

How do we know about it?

Speaker: Srini Kompella

39:17
Yeah, Very, very similar thoughts.
I mean, you know, so I probably, you know, encapsulate all of it into, you know, almost like a, you know, like a slogan that I believe the companies will need to drive.

AI is the new operating system for the business, right?
For pretty much every business.
I mean, that's powered with, you know, reliable, trusted data and it's, you know, rolled out, you know, for business value by the people.
But, you know, companies need to recognize that AI is fast becoming the operating system for, you know, the businesses that then starts the entire, you know, cascading down of, you know, using AI and everything that people are doing, right?
You know, obviously it has to drive business value. So far, for example, we talked about driving business value with, you know, data analytics and AI.

Let's not forget, I mean, from HCLTech. We also have a compelling proposition, you know, a platform proposition called AI Force that focuses on applying AI to do, you know, your software development life cycle better. To do your IT operations better, to do your business process operations better.

So, so that's the, you know, the efficiency piece, right? So which is equally important to the effectiveness piece with, you know, business value with the data and AI.

So, but for all of this to happen, you know, companies need to literally look at it as the AI is the new operating system going forward, which means that, you know, you got to look at how do you rewire, you know, your operating model.

Jeeva talked about it and you know, Andy, as you know, in HCLTech, we talk a lot about the operating model shifts that need to happen because of AI and with AI, right, then the people, you know, side of it, which is, you know, it's easy enough to say that, you know, I like what Jeeva says, right?
Don't leave the people behind.
It's easy enough to say that, hey, I mean, we need to get new AI to do things.
But you know, there is a rich levels of experience and context that's of the people.
How can enterprises tap into that and augment AI to uplift and then the processes, right?

How do you rewire the processes?

And, you know, now we see enough, you know, especially on the consumer AI, you know, you see enough, you know, entrepreneurial spirit coming in on monetizing, you know, with AI.

So that then is the, you know, the real opportunity that most companies have.

How do you create new revenue streams with AI? But for all of this, you know, I think the one line that I would say is that, you know, can companies start looking at AI as the new operating system of the future?

Speaker: Jeeva AKR
41:42
Beautiful.
Yeah.

Moderator: Dr. Andy Packham
41:44
Brilliant.
Thank you. I think look, yeah. For, me, my heart is firmly in operations.
That's my background.

When I start hearing, I know it's sad.
When I start hearing about BCP and DR and SLA as and operations, you know, that's music to my ears.
That's when it's really starting to get real.
It's not the sort of what we could do this or we could do that, but you start talking about DR and 3:00 in the morning calls, you know, this is recently real scale.

So yeah, I think this has been a really insightful conversation.
I think we've kind of learned a lot about how scaling is not about the tech.

It's about the discipline, it's about the governance, it's about the cost, it's about the people and how, you know, HCLTechs, AI Foundry and Azure just worked together very tightly integrated to deliver that and to deliver that reliably.

So Jeeva, Shrini, thank you very much for being a brilliant conversation.
I've really appreciated it.

Thank you.

Speaker: Jeeva AKR

42:44
Thank you.

Speaker: Srini Kompella

42:45
Yeah, thank you. It's been an absolute pleasure.